DOCUMENT RESUME

ED 164 010 IR 006 829

AUTHOR Gammill, Robert C.
TITLE Personal Computers for Science in the 1980's.
REPORT NO P-5954
PUB DATE Feb 78
NOTE 27p.; Paper presented at the Conference on Problems of the 80's: The Oregon Report on Computing (Portland, Oregon, March 21, 1978)
AVAILABLE FROM Rand Corporation, Santa Monica, California 90406 ($1.50)

EDRS PRICE MF-$0.83 Plus Postage. HC Not Available from EDRS.
DESCRIPTORS *Computers; Economic Factors; *Equipment Standards; *Futures (of Society); *Performance Specifications; *Scientific Research; Technological Advancement
IDENTIFIERS *Personal Computers

ABSTRACT

Increased use of personal computers for scientific research could reduce the dependence of scientists upon organizational affiliations, geographic location, and funding agencies, thereby improving the capabilities of isolated institutions of learning and research, and making rural localities more amenable to high technology. The scientific research base could also be increased as the tools of research support become more widely available. The achievement of these goals includes a number of technical and economic problems, primary of which are secondary storage, display, operating system design, device reliability, and inter-computer communications. The critical long-term problem area, where slow progress is likely, is economical inter-computer communications. The personal computer is more independent, flexible, and dependable and is capable of providing the following general functions because it is not shared: (1) tight coupling between the user or laboratory equipment and programs, (2) dependable access to archival information, (3) user tailoring of the terminal interface and services provided with minimal constraints, (4) user control, (5) continuity of service, (6) independence, and (7) wide access to other computer services. (CWM)

****
* Reproductions supplied by EDRS are the best that can be made *
* from the original document. *
****

ED164010

U S DEPARTMENT OF HEALTH,
EDUCATION & WELFARE
NATIONAL INSTITUTE OF
EDUCATION

THIS DOCUMENT HAS BEEN REPRODUCED EXACTLY AS RECEIVED FROM THE PERSON OR ORGANIZATION ORIGINATING IT. POINTS OF VIEW OR OPINIONS STATED DO NOT NECESSARILY REPRESENT OFFICIAL NATIONAL INSTITUTE OF EDUCATION POSITION OR POLICY

PERSONAL COMPUTERS FOR SCIENCE IN THE 1980's

ROBERT C. GAMMILL

FEBRUARY 1978

"PERMISSION TO REPRODUCE THIS MATERIAL IN MICROFICHE ONLY HAS BEEN GRANTED BY

Rand Corp.

TO THE EDUCATIONAL RESOURCES INFORMATION CENTER (ERIC) AND USERS OF THE ERIC SYSTEM."

"Problems of the 80's: the Oregon Report on Computing" --
Invited paper presented at the Session on Personal Computing, March 21, 1978, Portland, Oregon, co-sponsored by the Computer Science Departments of Oregon State University and the University of Southwestern Louisiana with the cooperation of the ACM and IEEE Computer Society.

IR006829

Personal Computers for Science in the 1980's

by

Dr. Robert C. Gammill

ABSTRACT

This paper examines the issues stemming from the projected use of personal computers (PC) in support of scientific research in the 1980's. Although personal computers will certainly have much wider application, the paper concentrates on this one area in order to reduce the effect of numerous non-technical factors which will dominate the growth patterns of personal computing in the wider arena, especially home and recreation uses. Business applications are also ignored, not because of non-technical issues but because specialization seems required. Undoubtedly some insights will prove transferable. The paper attempts to identify the issues, by carefully describing our model of the PC, identifying its advantages over present shared scientific computational facilities, assessing the problems imposed by inter-computer communications, and examining present and near-future hardware and software capabilities and constraints in light of the tight cost limits on personal computers. A conclusion reached is that the availability of PC's could lead to significantly increased capabilities for many research scientists, and reduce the dependence of scientists upon organizational affiliations, geographic location and funding agencies. However, the achievement of this goal includes a number of technical and economic problems. Primary among these appear to be secondary storage, display, operating system design, device reliability and inter-computer communications. The important problem for the early 1980's seems to be secondary memory. The critical long-term problem area, where slow progress is likely, is economical inter-computer communications.

INTRODUCTION

This paper examines the problems involved in the development of an inexpensive personal computer for use by scientific researchers in a variety of disciplines. Since the term "personal computer" (PC) is ill defined, we begin by describing what it is intended to mean here, and the special requirements imposed by the context of use in scientific research. Subsequently there will be much more detailed examination of the reasons behind the assertions that follow.

To begin, it should be pointed out that we will be using the term "personal computer" in a somewhat different manner than is common among those involved in hobby computing. This will certainly cause problems for some readers, but no obvious solution presents itself. We urge the reader not to forget that the context of this discussion of personal computation is that of support of scientific research, and that on some occasions this may cause conclusions to be reached which are less reasonable in the context of hobby computing. Briefly, it is our view that "private" is a synonym for "personal" when speaking of the personal computer for scientific research. It is visualized as an unshared computer facility for carrying out time-critical functions where convenience, guaranteed access, guaranteed response or other such issues are more important than the actual unit costs of computer resources such as processing and on-line storage. Although it is probable that the personal computer for scientific research would find application in wider areas, we do not feel that it is useful (in the context of application to scientific research) to use "hobby", "recreational", "small business", "home", "tiny" or even "micro" as synonyms for "personal". It is undeniable that these adjectives have some applicability.

A primary requirement of a personal computer for scientific research is that it be a dependable tool for generating, manipulating and examining text and data collections of limited size. We assume this to include some limited numerical computation. A second requirement is to serve as an interface to digital communication utilities, quickly interchanging data and text with research peers, administrative supervisors, publication editors and remote computers that have resources or services not available in the personal computer. A third requirement is to serve as a dependable and unsleeping agent, often in the user's absence, carrying out requested tasks and collecting incoming data and messages. These requirements demand that the personal computer be a general purpose system, with a flexible multi-processing operating system, and substantial but inexpensive demountable file storage for archiving. Rapid advances in LSI

electronics technology promise the availability of marginally suitable hardware configurations at an appropriate price ($1000 to $4000) in the next twelve years. In addition, the software necessary to complete the personal computer for scientific research presents a substantial set of problems. The problems stem from the small size and resource limits imposed by low cost, and the elimination of sharing (the method of the 70's) as a means of lowering costs. It is the interaction between these cost and resource limits, and the requirements for very powerful and flexible function, that provide the research problems of personal computers. Since the hardware limits of concern are in the future, a part of this paper will attempt to predict the trends of micro-computer technology and economics in the 1980's.

PC problems for the 1980's fall in several areas of concern:

a) Human interface - the means by which the user and computer communicate. This includes graphics display, user aids, text editors, discipline tailored application packages, and other aspects of the user environment including convenience of offline storage. A major problem here is that many scientists are computer naive (although many are better prepared than the general public).

b) Communications - communicating with remote terminals or computers. The rapid growth of digital communications, e.g., the ARPANET, value added nets and dial-up digital communication over telephones, necessitates software and hardware for dealing with increasingly complex communication problems. In addition, few standards exist here.

c) Operating System - the demands of the two preceding areas, along with the necessity for good dependability, require an operating system of a sophisticated and flexible nature in a machine of quite limited resources.

d) Secondary Memory - a short term hardware problem stemming from the need for substantial capacities, data rates and reliability at very low costs. The need for elimination of electro-mechanical components, because of their generally lower reliability, is also a part of this problem.

e) Reliability and maintenance - the PC will be one of the most complex consumer products available. Since we expect the PC to become a critical personal tool, dependability becomes an essential precondition. Increasing circuit integration should mitigate this problem, but is unlikely to eliminate it.

Research related to solving certain problems in these areas has been carried out at The Rand Corporation, using DEC PDP11 computers running the Bell Laboratories UNIX Operating System (Thompson 74). Although this is much too bulky and expensive a system to be called a PC, it is our expectation that it will be possible to purchase a PC analog to this system in the 1980's. Work at Bell Labs (Lycklama 77) using the DEC LSI11 computer is attempting to move toward that goal at the present time. Work at Rand has concentrated in the areas of intelligent agents, packet switched network interfaces and protocols, dial-up interfaces, mini-computer operating systems, user oriented tools to remind (send a message or run a program at a specified time), two dimensional CRT text editors (easier for naive users to learn and use than traditional line editors) as described by Bilofsky 77, and CRT multi-window user control of parallel processes.

WHY PERSONAL COMPUTERS?

A number of papers and articles have been written in recent years describing the promise of the personal computer, including Horn 75, Curran 77, Gumpert 77, Isaacson 77, Johnston 77, Kay 77, and Warren 77 (1 and 2). A basic question that arises in any technical article is, what advantage the personal computer provides over terminal access to a shared computer facility. The answers are dependability, independence and flexibility. It should be emphasized that these advantages are in addition to the obvious advantages of sharing, which have been carefully included in our model of the personal computer for scientific research (PCSR would be precise, but PC will be used for brevity). Those answers would appear to be contradicted by the limited resources that will be available in a personal computer of moderate to low cost. We examine each of these characteristics below, and show how the PC can provide each:

a) Dependability - the PC will serve in two major modes, stand-alone and as a front-end to other systems. It is due to the ability of the PC to serve in front-end mode that none of the good features of terminal access to shared systems need be lost. However, we assume that time-critical administrative tasks such as text editing, filing, work schedule maintenance and reminders will generally be carried out in stand-alone mode. Thus, when using remote systems with the PC serving as front-end, should any component (other than the PC itself) fail, the user can concentrate on stand-alone administrative tasks until service is restored. This does demand that the PC be an extremely reliable piece of equipment, but this seems technically feasible for future LSI electronics.

This demand lends urgency to the need for eliminating all electro-mechanical components, i.e., magnetically recorded mass memory such as tapes or disks. Another form of dependability is that denial of service (e.g., via a change in the user's job, junking of an old processor, overcharging, or scheduling of preventive maintenance) no longer can make access to archived files difficult or render a critical piece of software useless.

b) Independence - in all shared systems, the community of users and the economic and usage policies imposed by the system management and system architecture constrain the individual user to "stay with the crowd." The constraints differ for different systems, but the user often finds that what he desires cannot reasonably be accomplished in the community machine. A classic example of such an application is the real-time control of laboratory experiments. Ownership of a PC allows many more alternatives (within its resource limits) since only the owner need be satisfied with that part of the solution implemented in the PC. Also, the costs will usually be limited to the owner's time to develop the PC software, since the PC is expected to be self maintaining, i.e. sufficient for developing and debugging all its own software.

c) Flexibility - the trade-off between stand-alone and front-end modes allows great flexibility in choosing how to implement applications. Also the ability to tailor the PC precisely to the owner's needs adds a dimension of flexibility. Most important is that the PC provides a tool to aid the user in coping with the computational market place. When shared resources are available at a good price and appropriate form, the PC owner can use the PC as a front-end aid and dependability enhancer. On the other hand, in a seller's market, the user has improved alternatives such as stand-alone use or submission of tasks to very-remote computers via night-time long distance communication when the user is absent.

The analogy of a personal automobile in conjunction with mass transit, airlines, railroads and other transport services, has some usefulness in understanding the motivation towards personal computers. In the center of some large cities there is no need for a personal automobile (unless the transit workers strike), but most places the personal automobile is absolutely essential as a supplement to the limited public transport facilities provided. Some places are so isolated that the personal automobile is the only feasible solution. Note that "public"

usually means shared, and "personal" normally means unshared. The personal automobile is a tool that allows the individual to make cost and convenience trade-offs. In the transportation field, the Volkswagen beetle provided an inexpensive tool for coping. The extent to which that is important can be seen in the number of Volkswagens used for cross-country travel, although clearly not designed for that use. The personal computer offers similar kinds of opportunities. However, the operatorless (agent) aspect of the PC does not fit the transportation analogy. An agent is like an unsleeping personal secretary, able to accept incoming correspondence and issue outgoing correspondence, send reminders or give commands at times specified by the boss. Of course, these agent facilities can be provided in shared computers, but the implicit privacy of the PC (providable by software in shared systems) and the dependability make such services especially attractive in an owned PC.

In conjunction with the discussion of the preceding analogies, it may be useful to examine one further analogy for the purpose of clarifying what is meant by the adjective "personal". The personal automobile is frequently shared, as parents with teenage children can testify. In such situations, "family" is more appropriate than "personal". However, the desk that many employees use, though rarely owned by its user, is almost never shared. This is because so many things of personal importance are stored there. The user of a desk does not expect that others will rummage through that desk, especially not the employer. In other words, there is a social contract stating that repositories of personal information and effects become "personal." If the employer wishes to reclaim the desk, it is expected that some warning will be given, so that personal effects may be removed. In other words, although the desk is serially reusable, its reuse has a time scale of about a day. It is this lack of sharing, and exercise of personal control, that is the center of our use of the word "personal" in personal computers. It is interesting to note, in addition, that the potential for personalization of a computer system is much greater than that of a desk, so in the future employers and employees may come into conflict over ownership, allowable customization and access rights to personal computers used in an employment situation. Thus, it may become necessary for employees to purchase personal computers for use at work, in order to protect their rights to access control.

Another question concerning personal computers is, "What problem does the personal computer solve?" The question is normally asked by a computer user who is surrounded by a wide variety of shared computer facilities, and has few limits on the funds that can be spent for purchasing computer services. The answer is, better distribution of computer services to small organizations,

localities and to individuals. This is in direct analogy to the impact of the automobile on those living in areas where convenient public transportation is not feasible. At present, access to computer services tends to be concentrated in the hands of the big, the rich and the centralized. Thus, in scientific research, the small college or university labors at a disadvantage in attracting and keeping researchers and grants which depend upon good computational support. This problem compounds the difficulties of rural areas in attempting to attract high technology industries which desire educational support in computer related subjects. Thus, the personal computer, like the inexpensive personal automobile and the TV set, could serve as a tool for revitalizing rural areas.

WHAT IS A PERSONAL COMPUTER?

There are many possible versions of what a personal computer should be. The goal of the following discussion is to distinguish our model of the PC from a number of other models and to illustrate why the distinctions are important in defining the functional capabilities that the PC can provide. Then we will examine the technological limits imposed by present and projected hardware and the software research questions which result from the interaction of these with the functional requirements of the personal computer. We will proceed by describing several classes of computer systems on the present market which might be mistaken for personal computers, and then contrast the PC (which has no present realization at any suitable price) with those. To summarize what will be said, the PC is a general purpose multi-processing computer system, able to interface to a variety of external long-distance public communications media and able to run substantial user-developed software subsystems. The PC will have a minimum of mechanical subsystems, in order to achieve low cost, high dependability, high speed, low energy and low noise operation. A critical constraint is that the PC must be affordable by most scientists from their own personal resources. This constraint might seem unnecessary for institutionally supported research, but we contend that many researchers will be unwilling to commit themselves to dependence upon a critical tool, such as the personal computer, if continuing access to that tool depends upon maintaining a relationship with a specific agency or institution. Now let us proceed to examine a taxonomy of present small computer systems which can be mistaken for the personal computer. It should be noted that each of the categories has fuzzy boundaries, and that a series of changes to a system that falls in one of the categories may make it different in emphasis, moving it into another category. The categories are presented here only as a means for increasing the clarity of understanding, not as a firm taxonomy for zealous

defense.

Computer Terminal with a Micro-Processor (Intelligent Terminal)

Computer terminal manufacturers have discovered that adding a micro processor can make a terminal much easier to reconfigure. Well over 100 commercial computer terminals now incorporate a processor. Most of these are not user programmable, nor do most have any form of local secondary storage. Those that do have these features usually do not support the functions that are normally associated with a general purpose computer system, such as operating system, compilers, and problem oriented software. These terminals usually can interface to external communications, although the protocols they implement are very limited.

Stored Program Calculators

A number of hand calculators now have stored programs, which can be read from and written to magnetic storage cards. These machines are not general purpose, since they implement calculations in a user oriented machine language, and do not interface to external communications lines. However, the most important flaw is the inability to write an operating system.

Applications Oriented Computers

These are close to being PC's, and may be hard to distinguish from the PC. These computers allow scientific problem solving in a programming language such as Basic or APL. Examples are the WANG, Tektronix and IBM 5100 machines. The important distinctions are the lack of external communication (other than to the user or local mass store), the lack of a general multi-processing operating system, and dedication to a user oriented numerical problem solving language unsuitable for systems implementation. The latter means that these systems are not self maintaining, i.e. cannot be used to extend or modify their own system software, which is a significant deficiency.

Limited General Purpose Machines (micro-processors)

These are the micro-processors which are appearing on the hobbyist market. Issues of concern in this environment are examined by Allison 77, Garland 77, Mooers 77 and Warren 77 (1 and 2). These systems are normally self maintaining (in assembly language), but lack the sophisticated operating system and external inter-computer communications necessary for the PC. Examples of systems that are close to the desired price (at present), but too limited in hardware function, include the MITS ALTAIR and IMSAI computers, based on the INTEL 8080 micro

processor. These are 8 bit microprocessors, with a rather limited instruction set. It is the introduction of microprocessors which provide the capabilities of a 16 bit minicomputer that promises to support the personal computer systems of the 1980's. Examples are the DEC LSI11 and the micro NOVA. A complete PC system (under our definition) is still too expensive, primarily because of the high cost of peripheral devices, but prices are decreasing. Despite the appearance of much new hardware that could soon satisfy the cost constraints imposed on the PC, a big limitation is the lack of suitable software. Most important in this lack is software to allow inter-computer communication and the necessary operating system capabilities to support that.

General Purpose CPU with Inter-Computer Communications (our PC)

Although a stand-alone general-purpose machine satisfies many hobbyists as being the personal computer, our experience leads us to assert that unless it is interfaced to public digital communication utilities (e.g., telephone and/or commercial packet switched networks), in order to communicate with remote machines, the PC cannot provide the needed critical functions. In particular, when tasks too large for its limited resources are encountered, the stand-alone machine becomes an extraneous piece of gear, not intimately involved in the problem solving process. This means that the user will see a definite barrier at problem sizes that are too big for the stand-alone machine to handle. Thus, extensive labor that may have been expended toward making the computer useful for smaller problems in the area will be wasted, or at least hard to utilize. This kind of inflexibility is unacceptable in a critical personal tool. The interdependence of the computation and telecommunications environments has been explored by Anderson 77, Cerf 77, Farber 77, Greenberger 74, Rosenthal 76 and Wyatt 76. It should be emphasized that the requirement to serve as a gateway to other computing resources provides a strong set of constraints upon the architecture of the computer system. It is these constraints that radically alter the character of the software problems involved in the PC from those encountered in systems designed primarily for stand-alone use. Some of the specialized issues involved are examined by Akkoyunlu 72.

An example of a system that fits many of the requirements for a PC at present is the Digital Equipment Corporation PDP11 computer, running the Bell Labs UNIX operating system, described by Thompson 74, and later by Lycklama 77. Prices are still too high. A PDP11/34 costs around $30,000 or an LSI11 costs around $10,000, and these prices ignore the cost of the software (about the same). This system has been interfaced to packet switched

digital communication networks, and provides multi-processing and other operating systems functions, as well as allowing user modifications to the system and development of subsystems via a multitude of software development tools.

It will be our purpose in the material that follows to expand on the concepts presented above. It should be noted that we will be deemphasizing the provision of tools for numerical problem solving, although that is undoubtedly an important function of any computer to be used by a scientific researcher. Our reasoning is that all general purpose computers to date have included numerous operations for carrying out mathematics. The result is that most general purpose computers can be used quite effectively for numerical problem solving. However, computers and systems designed specifically for this purpose (such as Basic and APL systems) tend to be very rigid when it comes to being able to provide operating system and communications extensions. Thus, it will be our contention that the PC will be an easy system in which to write APL and Basic interpreters, or other numerical computation subsystems, and that these areas need receive no further attention here. It is expected that early users of the PC will supply such subsystems, and this expectation is supported by the history of many systems, e.g., UNIX or TENEX, described respectively by Thompson 77 and Bobrow 72. It is our contention, however, that lacks in the architecture, operating system or communications interface designs cannot be remedied so easily. Therefore, we will be concentrating upon these more fundamental elements in our study of the PC for scientific research.

WHAT FUNCTIONS SHOULD THE PC SUPPLY?

First, it must be pointed out that this is the most basic research question involved in the study of personal computers. It is not obvious what functions should be supplied. For those who have used a general purpose time-shared operating system (such as TENEX or UNIX), interfaced to a packet-switched communications network (such as the ARPANET), and providing substantial user software tools for generating application subsystems and manipulating and sending text, there exists some consensus about an intersection of critical functions that must be supplied. However, the set of functions about which that consensus exists is small. It should include a hierarchical file and directory system (TENEX users may not desire hierarchy, not being used to it), a good text editor, compilers and assemblers capable of maintaining the system's own code (to allow extension), an external communications interface, and an operating system capable of supporting all of this. As one examines a greater range of function, disagreement becomes

rampant. In this section we will examine a set of functions that we view as critical, as a result of experience with UNIX, and discussion with others. We note again that this is an area in which strong disagreement is guaranteed. Some of the disagreement stems from a commitment to sharing of resources, some from attitudes about what kinds of functions small computers can do well, and some from basic biases caused by the kind of computer service environment that an individual has experienced previously. The PC we envision must be able to serve as a stand-alone machine, as a user terminal simulator to remote machines, and at all the intermediate points of that spectrum, since we view the economic limits placed upon the PC as too severe to allow general scientific research to be carried out solely with its use. At a minimum, the PC should be able to serve scientists who write large FORTRAN programs as a dramatic improvement over the keypunch and decks of punched cards. For this to work, prepared source texts must be easily transmitted to the large central computer, and outputs generated by that remote computer must be easily collected and examined through use of the personal computer.

There are a number of viewpoints available concerning functions of the PC. First there is a general view, which deals with the functions that are available as a result of the fact that a personal computer is dedicated to the needs of just one individual. These functions are those general services which can be supplied by a personal computer because it is not shared with others. Despite the many advantages of sharing and centralization of function, there are functionalities which are eliminated or degraded by them. It is an implicit assumption in all that is said here that the personal computer will be useable in both stand-alone and front-end modes, so that optimization of function may be manipulated toward or away from sharing of computer resources, depending upon the situation. The following are some general functions which are available because the personal computer is not a shared piece of equipment.

a) Tight coupling (high bandwidth) between the user and programs or between laboratory equipment and programs. For example, this can allow the use of highly interactive text editors, where a single keystroke may evoke a substantial processing, or the real-time control of an experimental device, where a response time must be guaranteed. The tighter the coupling needed, the less feasible a shared resource environment is for supporting it.

b) Dependable access to archival information. "Archival" implies that there is a great deal of information with

low likelihood that a particular part will be accessed. Denial of access to the archived information can be a critical problem at certain times. In a shared environment, control over access resides with the management, not the user. The function of interfacing the user's archives to other systems can make the PC serve as a portable user-owned peripheral device. This provides insurance against media obsolescence, e.g., having data on cards with card readers becoming uncommon.

c) User tailoring (personalization) of the terminal interface and services provided with minimal constraints. The user often has unique applications and idiosyncratic usage patterns that can make the use of a facility aimed at the average user unpleasant or intolerable. Unique requirements may include the need for a special input device such as a special keyboard or graphic input device e.g., a map digitizer. Note that when the user can afford the computer, the incentive for training himself in the use of a special input device (e.g., stenotype) may become higher, since the utility of that training is less likely to be subsequently nullified.

d) User control (self determination). The user sometimes needs functions that are difficult or uneconomic to supply in shared facilities. Also, since the priorities of service and resource provision in shared facilities are controlled by economics and the management, the user may find certain services will not be made available at an acceptable cost (dollars, time, effort, etc.).

e) Continuity of service. The user sometimes needs to be able to maintain the hardware environment of critical programs e.g., running an old system after others lose interest, or to guarantee service at specific times of day, e.g., supply a file transfer service during a block of allocated time on a specific remote computer.

f) Independence. The user of computer services is sometimes "isolated", either geographically or in the sense of requiring unique services, such that the economies of resource sharing are difficult to realize. If sharing is the only feasible means for providing the service at an acceptable cost, either the isolation must be reduced or the service cannot be provided. The personal computer should ease some of this cost bias against isolation.

g) Wide access to other computer services. The present computer user has little control over the interfacing

characteristics of the local host computer (or other communication protocol mechanisms). As a result, access to remote computer facilities depends upon what local interfaces have been provided. These may not suffice for the individual user's needs.

Besides the general functional characteristics of a personal computer, as compared to a shared facility, there are the specific functions that should be provided by a general purpose personal system. The problem is that the specific type of function a user desires tends to depend a great deal upon what kind of service has been encountered in the past. Furthermore, as was noted earlier, certain functions are easily supplied by the user community, if more critical supporting functions relating to the operating system and communications interfacing are sufficiently powerful and flexible. Thus, we shall now examine those critical user-related functions that tend to be independent of the application area involved.

Most critical to the user is the provision of a very flexible file and directory system, allowing the user to structure collections of information in a manner appropriate to its use. Especially important is the capability to mount and dismount storage volumes, allowing access to an amount of information controlled only by the cost of the media to store it. It should be noted that some of the best operating systems, e.g., TENEX, do not allow routine removal of volumes. Despite the fact that UNIX does allow removal of volumes, it was not designed with routine mounting and dismounting of volumes by users in mind, and makes no attempt to keep an inventory of the contents of off-line volumes. This could be an important issue in the PC. Coupled very closely with the file and directory system is the text editing facilities by which new files may be generated, and existing files examined and modified. Two general kinds of text editors are essential to the easy use of a personal system. These are interactive and systematic editors. The function of the interactive editor is to provide very tightly coupled intercommunication with the user, allowing rapid visual scanning of text, and modification operations such as marking out, inserting and appending. Such editors often display text in two dimensions (not just as lines) and provide complex operations which can be specified by a single key-stroke, Bilofsky 77. Lay-up, or cut and paste type operations, tend to be easy to do with these editors. These tend to be most valuable for writing documents, laying out data tables and typing diagrams. Systematic editors concentrate more on global text substitution and modification operations, treating a file as a linear text string. These editors tend to be most valuable in modifying text which contains systematic errors. Both sorts of editing capability are critical

for a personal system in which a scientific researcher will carry out the majority of his administrative and clerical activities.

Communication, encompassing remote terminal simulation by the personal computer, interchange of files with remote systems, and perhaps remote job-entry, is also a critical function of the personal computer system. These capabilities, especially in conjunction with the ability to carry them out with the user absent, can allow the personal system to serve as an agent, controlling, prodding, querying, retrieving, and otherwise negotiating with remote systems for the benefit of the user. These capabilities are essential not only for communicating with remote computers, but for transmitting and receiving information with other persons. Examples of persons that a scientific researcher needs to communicate with include research peers, granting agency personnel, journal editors and librarians of specialized information collections. A special advantage of the personal computer in such communication is that sending of notes or documents to a distribution list is no more difficult for the user than sending to a single recipient. Notice that PC mail provides the advantage that the recipient need not be present to receive the mail (just as with present mail), but distribution is at electronic speeds. Another advantage is that forwarding can be carried out automatically, and remote mailing lists can take care of distribution to recipients whose names need not be known to the sender. The highest probable payoff from the provision of a personal computer system for scientific researchers would stem from the conjoining of the inter-personal communications system (typewriters, secretaries and mail) with the inter-computer communication system (cards, tapes and digital communications), in the form of a single device able to do both, and providing improved capabilities for dealing with the information under the control of the user.

Another critical function for the user, is multi-processing (meaning multiple processes in concurrent execution). Even though the user is immersed in a text editing or other time critical task, it is essential that the computer be able to continue acting as an agent, receiving incoming mail, negotiating with remote systems, or perhaps controlling a nearby experiment. The amount of multi-processing needed may be low, but the user will not be able to give up administrative tasks on days when real-time control of an experiment is required. It should be emphasized that multi-processing in the PC need not be implemented with multiple processors, nor is the apparently attractive option of a multi-processor PC excluded.

Another critical function is for the personal computer system to be self maintaining. This means that the user of the system

could modify its source code and compile a new version when that is desired. Since many of its users will want to develop subsystems that are relevant to a specific scientific research area, and many of those will have no other convenient software development system available to them, it is essential that the personal computer for scientific research be able to modify and extend, as well as debug, its own software.

When particular areas of application of the personal computer in scientific research are examined, many further areas of critical function can be identified. However, in this paper we place these in secondary importance to those requirements which cut across all application areas. As examples, for research in fluid dynamics, provision of substantial amounts of mass storage and graphic display capabilities will almost certainly be critical. For research requiring rapid interaction with experimental apparatus, high speed communication interfaces will also be a critical need. However, we believe that the identification of application dependent needs for critical functions must be the subject of future PC research.

OBSTACLES: PROJECTED TECHNICAL AND COST PROBLEMS

Before we proceed, it must be asserted that this is another area requiring fundamental research in personal computing. Some facts are known, although a careful survey and gathering of supporting evidence needs to be carried out. The dramatic and continuing decrease in cost and size, and the increase of function provided by large scale integration of solid state devices promises to bring the processing unit and primary memory of a personal computer well within range in the next couple of years. The problems generally reside in those areas requiring mechanical movement, high voltages or other complexity due to interfacing with the outside world. The specific problem components appear to be the secondary store, communications interfacing, and display. We shall touch on each of these briefly.

The secondary memory appears to be the critical near-term hardware problem for the personal computer for scientific research. Although this is a problem area for shared computer systems too, the technique that has been used for solving it is cost sharing among the users, which is not available in personal computing. An additional cost problem for unshared secondary memory stems from the fact that secondary memory can be viewed as having two components of cost. The first is the cost of the basic device, assuming zero storage capacity (heads, power supply, motor, controller, etc.) and the second is the cost of the storage. In the smaller secondary memories, the first cost completely dominates the second. In addition, the cost per unit

of storage space decreases as the total amount of storage increases, Brechtlein 77. This is "economy of scale," a cost reduction that is also not available to personal computers. In the large secondary memories used by shared systems, the cost per unit storage is viewed as primary and the basic cost thrown in as an overhead item. Somewhat the reverse view of things must be taken in personal computing. Ultimately this should lead toward the development of simpler (and cheaper) basic devices, so that there can be less economic emphasis upon large size in secondary memory.

The best candidates for PC secondary memory, in the near future, appear to be the floppy disk and the small hard-disk drive. The floppy disk has limited promise for improved capability, presently storing around two to eight million bits per drive. Advantages are the ease of use and low cost, which should come a little lower. However, the slow speeds required by head-contact with the media surface, and the wear that too much use entails, preclude the use of this device as a means for swapping storage images in multi-processing. The special bonus of the floppy is, of course, the easy user handling of a high-density random access medium. The small hard-disk drive doesn't have the head contact problem, but costs are too high and have limited promise for coming down much in the future, and user dismounting and storage of these volumes is awkward. Providing some hope is a recent note (Electronics, Jan. 19, 1978, page 33) stating that Shugart, recently purchased by Xerox, and IBM are working on Winchester related hard disks (not demountable) to provide 10 to 100 megabytes for less than $2000. This would be an important development! At present (Jan. 78) Ohio Scientific advertizes a 74 megabyte Winchester disk with controller for $6,000. Since these hard-disks are usually not demountable, they would normally be used in addition to floppies, so cost is still a problem in the 1980's. Some new technologies such as charge-coupled-devices (CCD) give some promise of helping to solve the secondary memory problem, by serving as swapping memory, perhaps in conjunction with floppy disks for user files. Bubble memory devices (BMD) may be able to supplant the floppy disk in the 1980's, but progress in cost, capacity and speed will be necessary. Another possible solution is that RAM memory may reach such low prices that intermediate secondary-memory for swapping may not be needed, and multi- processing may be carried out through memory mapping. All of these alternatives need examination. The one uniformity across all of these possible scenarios for secondary memory is that the operating system and communication modules will require extra sophistication for coping with the special problems imposed by inexpensive and limited secondary-memory hardware, especially any hybrid approach. Most desirable for the software of a personal computer for scientific research would be

an operating system designed to deal with a variety of secondary storage mechanisms, including varying sizes, speeds, and types of memory. Clearly secondary memory (at the level of function we demand) will be a cost problem throughout the 1980's, and only an unexpectedly rapid decrease in cost by one of the technologies could eliminate this as a PC problem area. Another reason why secondary memory is of concern is that it is the major remaining electro-mechanical component, and a primary site of potential unreliability. However, despite some pessimism, steady progress can be shown in the area of secondary memory, so despite the fact that it may still be a problem in the early 1980's, it appears likely that secondary memory will cease to be a problem towards the end of the 1980's.

With respect to communications, the problem is not only a lack of sufficiently powerful equipment at acceptable prices, but a problem of a complex and unpredictable communications environment. The wrangling between the public, communications companies and monopolies, regulatory agencies and other government bodies, promises to keep communications an unpredictable and frustrating area, where much less will be accomplished than is technically feasible. The problem is documented by Oettinger 77, in the annual report of the Harvard Program on Information Policy. In addition, a personal computer will face a more complicated communications environment than do present day corporate or institutionally owned computers, since it will be expected to communicate with a wider and less predictable range of other systems. Another way of saying this is that present computer-computer communication tends to be homogeneous since both computers are usually owned by the same organization, and a pair of matched modems are purchased for their use. Even when the computers are incompatible, a homogeneous interface (e.g., the ARPA Net) is supplied and both must match themselves to that standard. However, the PC will often communicate with computers owned by strangers down lines not intended for that use (and in fact, likely to be conditioned in the hope of making such use impossible). As a result, a communciations interface for a personal computer for scientific research will have to deal with an arbitrary and capriciously changing environment. This will mean that its modems and associated interfaces will need to be very powerful and flexible (perhaps even adaptive) while staying inexpensive. In the short-term, this will probably include dial-in and dial-out interfaces, and acoustic and direct coupled modems for connecting to ordinary phone lines. However, in the long-term present devices and methods are likely to provide too limited a data rate. For example, if we assume 300 bps (or 30 characters per second) a 50 line by 60 character page (of 3000 characters) takes 100 seconds to send. Assuming $1 for night-time long distance 3

minute calls, a two page letter costs about a dollar to send. Approximately a factor of ten is needed for true feasibility. Admittedly, this analysis is oversimplified, since packet-switched networks can provide more economical high-speed long-haul communications, and 1200 bps full-duplex modems are available for dial-up lines (but for $1000+). Although many technical approaches to a solution can be hypothesized, including standards and commercial packet-switched networks, it appears likely that the multi-adversary nature of modern digital communications will make it a major stumbling block to effective use of personal computers throughout the 1980's and long after.

A third problem area is display. This is not a critical problem, as is secondary memory or communications, but is troublesome all the same. The display is the means by which the personal computer will communicate with the user, and just as a secretary is judged by the appearance of the final product, the relationship of the user to the computer will also be affected by the quality of presentation. At present inexpensive CRT displays are the primary display device, but these are bulky, fragile and require high voltages. Inexpensive CRT displays usually allow display of around 24 lines of 80 alphanumeric characters, but no graphics. Early versions of personal computers for scientific research will likely have a similar kind of display. Hard copy must be treated as an optional extra (since it is likely to remain quite expensive), and especially because communication capabilities could be used to transmit material to a remote shared printing service when hard copy is required. In the long term, display must move towards raster-scan graphics, and less bulky and fragile hardware. Plasma panels have many desirable characteristics for such use, but high cost and energy demands tend to put them out of the race. Numerous efforts are underway in Japan and the U.S., attempting to produce flat panel TV, but these are still laboratory efforts, so predictions of impact are almost impossible. Finally, it is also important that the display move in the direction of allowing a complete page of text (an eight and one half by eleven inch typed sheet) to be viewed at a time, further decreasing dependence upon hard copy. Forseeable development of CRT displays promise to produce such capabilities (Motorola has shown its M4408 CRT for doing exactly that at NCC77) but non-CRT competition appears very unpredictable. The Xerox Alto seems to have the most flexible display available at this time, but little is known about the cost and reliability of this device, since it has only recently moved toward commercial production. High resolution and accuracy coupled with low cost has proved an elusive goal in the display area.

Another aspect of the display problem is that the personal

computer enjoys a special advantage over shared computer systems in that its display can be viewed as an operator's console, instead of as a remote terminal. This means that direct-memory-access by the display, and modification of displayed text at memory speeds can easily be utilized to accomplish tight coupling to the user's inputs. However, this advantage brings with it some problems relating to increased loading of an already limited address space in most small machines, and an additional kind of device for which software drivers must be provided. Software solutions to these additional problems will be needed for the personal computer for scientific research. Again, the Xerox Alto (Kay 77) provides the best examples of use of these opportunities while dealing with the problems.

The preceding paragraphs should be viewed as rough examples of a kind of analysis which needs to be carried out in greater detail and depth by PC researchers. We conclude that a critical problem for the early 1980's will be the choices of primary and secondary memory that will be feasible for use by the personal computer, due to technical, convenience, reliability and cost constraints. We also conclude that inter-computer communication will be the problem that becomes of most critical importance as stand-alone machine users begin to demand it. We have stated considerable pessimism about progress in this area, both because of the number and sizes of political entities involved, and because past progress in this area has been relatively slow and unimaginative.

WHAT STEPS ARE NEEDED TO OVERCOME THE OBSTACLES?

In this section we will attempt to describe progress already underway toward solving the problems mentioned in the previous section, and certain areas where new initiatives must be undertaken.

The secondary memory problem is one in which steady (if slow) progress is being made. It is primarily a problem of mass production of inexpensive devices for which feasibility has been proved. As the PC market broadens, competitive forces should lead to suitable solutions to this problem. The exact timing of the solution (and its form) are difficult to predict, but this appears to be an area where PC researchers need not be concerned, and need only monitor progress in order to be prepared. The importance of increasing reliability (e.g., through elimination of electro-mechanical components, or sealing them inside a cover as in the Winchester disk) needs to be emphasized to the manufacturers of these devices. The development of ever more intelligent (micro-processor based) controllers seems inevitable, and should be encouraged as a method of achieving greater flexibility and reliability (e.g., error detection or correction)

while decreasing costs (due to LSI circuit cost reductions). The only possible reason for pessimism concerning secondary memory is the number of promising new technologies that have failed when introduced to the actual market place in the past. Optimism seems justified today, however, because the proven floppy and Winchester disk technologies provide fall-back (insurance) positions, although a little too expensive and unreliable.

In displays, upgrades of present CRT technologies using LSI circuitry appear to provide the most predictable path. Many other potential technologies are being developed in the laboratories, but the economics of competition with the well established CRT technology are complex. Many new technologies that could be used for computer display, will also be usable for television. Thus, most new display technologies have a large and well established market available when they reach readiness for production. However, competing against a technology as entrenched (and polished) as the CRT will require substantial advantages and low costs. The bottom line on display is that extrapolation of present CRT technology appears likely to satisfy the demands of a modest PC, so long as we keep our expectations low. However, the automated office desk and the DYNABOOK (Kay 77), demand much more than the CRT can supply.

Our model of the PC, which includes external communication to other computers, adds much complexity to the PC software, since multi-processing is required. In addition, as personal computers become more commonplace, their users will naturally become (on the average) less knowledgeable about computers. All of this will demand much improved software. However, unlike hardware research and development, software can be developed by almost anyone who owns a PC, and this includes members of the various departments of Computer Science throughout the nation. In addition, the relatively small size of the PC means that individuals and small teams can hope for some success in their efforts. Also, high quality systems implementation languages such as C, Smalltalk and Pascal have begun to appear on minicomputers, as has at least one good operating system (UNIX). All of this leads to some considerable optimism that, although software problems will continue to exist, the introduction of the PC with the broadened software development community (and market) that it creates, will certainly accelerate the quantity of new PC software (quality is another matter). The major problem (and reason for pessimism) with respect to PC software is that the PC makes software theft trivial, so the market for unbundled software will be extremely disorderly, discouraging commercial enterprises from seeking to enter that market. If there is any concern over PC software, it should be in finding a breakthrough in software protection mechanisms, so as to create an orderly and

unregulated market for quality PC software. The cost of such a breakthrough is difficult to predict, but its effect would clearly be to raise the quality and increase the amount of software available to the PC user.

Finally, let us examine the PC inter-computer communications problem. As has been mentioned a number of times previously, we believe this to be the critical long-term (towards the end of the 1980's and beyond) problem for PC. There are numerous components to the problem, many of them non-technical. In fact the most important components appear to be political. We will attempt to summarize the problem, but we believe that this area should receive much greater attention from serious researchers in PC. Part of the problem in "public" communication is that it has been primarily the province of relatively traditional electrical engineering technology. Incursions by modern computer science have been relatively few, most notable being the AT&T Digital Switching System and the ARPA Net. Another problem is that public digital communications has long been regulated by government (FCC etc.), and under the control of a monopoly, AT&T and its operating companies. The special problem is that this is the arena in which dominant organizations such as AT&T, IBM, Xerox, CBS, the U.S. Post Office, the FCC, and consumer organizations, will maneuver in the 1980's and beyond, with the outcome and winners uncertain (see Oettinger 77). In addition, most present computer communication is intra-company, and thus not subject to federal regulation, but PC communication will clearly be classified as public communication, and thus be at the center of the arena, and of immense potential value to the winners. Furthermore, the individual PC owner will be quite ineffective politically. Until PC oriented consumer organizations have sufficient political force to affect the arena, the only "defenders" of PC communications will be those who hope to gain power through it. In such an environment, pessimism and preparation for a siege would seem justified. In fact, that siege may already be starting, since new tariffs (e.g., metered service) being sought by the telephone companies appear to be in rather direct response to the use of telephone lines for digital communication. So, in this battlefield of giants, what hope is there for PC communications? The answer is, considerable hope, but only through use of technological leverage. Just as the duplicating machine in the hands of the public has had a dramatic impact on the publishing industry, the offspring of the acoustic coupled modem used as an integral part of a PC system could make capture of PC communications by the giants very difficult (if not impossible). The underlying reason is that personal voice communication (e.g., for the elderly to seek medical aid), although not written into the Bill of Rights, is a strong social contract between the U.S. government and the

voters. The removal of that right, or pricing it beyond the reach of many, would certainly produce a reaction sufficient to change the balance of power among the giants. Also, voice communication requires a bandwidth considerably greater (today) than that needed for PC communication. However, that bandwidth must be effectively used, and to date that has rarely been the case. In other words, PC communications must tie itself as closely as possible to the residential telephone user, in order to find protection. However, standing still will not be good enough, for the telephone companies will be seeking technological solutions for separating computer from voice communications, in order to charge what they consider a fair share for computer communications. Note that from the point of view of the phone companies, personal voice communication is under-charged, due to regulation of that market. As a result, modems for personal computers need to be developed which can use all of the channel capacity of lines which are intended for voice. It would also be very helpful if those modems could produce sounds that are relatively difficult to distinguish from voice through automatic means. It should be clear by now that what we are recommending for PC communications is akin to guerilla warfare, but on a technological battleground. No research relating to this kind of modem is known (by the writer) to be in progress at the present time. As a result, we feel that it is here that a very important breakthrough is required. The cost of that breakthrough is unknown. We believe it to be of great importance, since with this goal achieved, the PC in the hands of the public could become a very powerful tool.

One other point needs to be made in relation to PC communications. There exists an established (if wobbly) market for long-haul data communications at the present time. This is the market being served by Telenet Corporation. It is very possible that long-haul digital communications at reasonable rates may become sufficiently established so that the coming battles cannot eradicate it. However, the PC will still require communications from its location to the nearest host office or satellite dish. Thus, it appears that the most critical arena for PC communications is local. In addition, communications experts assert that the most arcane and difficult regulations are to be found at the local level. So, despite the fact that we cannot precisely predict the skirmishes nor winners of the coming battles over telecommunications, it is still clear that PC communications stands at a relatively weak position, is likely to be viewed as a rich prize for capture, and should begin to seek sophisticated defensive measures for protecting itself. As we have indicated, such measures seem feasible, but research seeking that defense needs to begin very soon.

SUMMARY

A particular model of a personal computer for scientific research has been described. The implications of that model have been investigated, both in respect to user function provided and problems of implementation. A number of research problems have been identified, which although they may bear resemblance to other computer science research problems, take a substantially different character in this context. A number of personal computing problem areas for the 1980's have been identified. These include: secondary storage, display, operating system design, reliability and inter-computer communications. It was stated that secondary storage will be the primary problem of the early 1980's, but that communications will be the long-term critical problem. It is felt that a breakthrough is needed in the latter area. It was concluded that the personal computer for scientific research could reduce the dependence of scientists upon organizational affiliations, geographic location and funding agencies. This, in turn, could improve the capabilities of isolated institutions of learning and research, making rural localities more amenable to high technology. This could also increase the scientific research base as the tools of research support become more widely available.

REFERENCES

E. Akkoyunlu, A. Bernstein, and R. Schantz, "An Operating System for a Network Environment," Proceedings of the Symposium on Computer-Communications Networks and Teletraffic, Polytechnic Inst. of Brooklyn, April 4-6, 1972, pp. 529-538.

D. R. Allison, guest ed., "Small Scale Computing - introduction", IEEE Computer, Vol. 10, No. 3, March 1977, p. 8.

R. H. Anderson and J. J. Gillogly, "The RAND Intelligent Terminal Agent (RITA) As A Network Access Aid," AFIPS Proc. Vol.45, National Computer Conf., New York, June 1976, pp. 501-509.

W. Bilofsky, "The CRT Text Editor NED - Introduction and Reference Manual," Rand report R-2176-ARPA, December, 1977.

D. G. Bobrow, J. D. Burchfiel, D. L. Murphy and R. S. Tomlinson, "TENEX, a paged time sharing system for the PDP-10," CACM, Vol. 15, No. 3, March 1972, pp 135-143.

R. Brechtlein, "Comparing Disc Technologies," Datamation, Jan. 1978, pp. 139-150.

V. G. Cerf and A. Curran, "The Future of Computer

Communications," Datamation, May 1977, pp. 105-114.

L. Curran, "Personal Computers Mean Business," Electronics, Vol. 50, No. 7, March 31, 1977, pp. 89-96.

D. Farber and P. Baran, "The Convergence of Computing and Telecommunications Systems," Science, Vol. 195, No. 4283, 18 March 1977, pp. 1166-1170.

H. Garland, "Design Innovations in Personal Computers," IEEE Computer, Vol. 10, No. 3, March 1977, pp. 24-27.

M. Greenberger et al. (eds), Networks for Research and Education, MIT Press, Cambridge, Mass., 1974.

D. Gumpert, "The Computer Moves From the Corporation To Your Living Room," Wall Street Journal, Vol. 96, No. 25, Feb. 4, 1977, page 1.

P. Isaacson, "Personal Computing Dissected," IEEE Computer, Vol. 10, No. 7, July 1977, pp. 71-73.

B. K. P. Horn and P. H. Winston, "Personal Computers," Datamation, Vol. 21, No. 5, May 1975, pp. 111-115.

T. Johnston, "Will Your Next Home Appliance be a Mini-Computer?", New West Magazine, March 14, 1977, pp. 50-59.

A. Kay and A. Goldberg, "Personal Dynamic Media," IEEE Computer, Vol. 10, No. 3, March 1977, pp. 31-41.

H. Lycklama, "UNIX on a Micro-processor," AFIPS Proc. Vol. 46, National Computer Conf., Dallas, June 1977, pp. 237-242.

A. Oettinger, Chmn., "Information Resources Policy: Arenas, Players and Stakes," Annual Report 1976-1977, Volume One, Program on Information Resources Policy, Harvard University, Cambridge, Mass. 02138, 1977.

R. Rosenthal, "Network Access Techniques--A Review," AFIPS Proc. Vol.45, National Computer Conf., New York, June 1976, pp 495-500.

J. Warren (1), "Personal and Hobby Computing: An Overview," IEEE Computer, Vol. 10, No. 3, March 1977, pp. 10-22.

J. Warren (2), "Personal Computing - An Overview for Computer Professionals," AFIPS Proc. Vol. 46, National Computer Conf., Dallas, June 1977, pp. 493-498.

K. Thompson and D. M. Ritchie, "The UNIX Time-Sharing System," CACM, Vol. 17, No. 7, July 1974, pp. 365-375.

J. B. Wyatt and V. I. Polley, "Network Interface Systems--An Evaluation By Simulation," AFIPS Proc. Vol.45, National Computer Conf., New York, June 1976, pp. 511-521.